

# Genetic Algorithm-based Mapper to Support Multiple Concurrent Users on Wireless Testbeds

Yaser A. Elnakieb  
Electrical and Computer Eng. Dept.  
Virginia Tech MENA  
Email: ynakieb@vt.edu

Michael Azmy  
Computer and Systems Eng. Dept.  
Faculty of Engineering,  
Alexandria University, Egypt  
Email: michael.azmy@alexu.edu.eg

Mustafa ElNainay  
Computer and Systems Eng. Dept.  
Faculty of Engineering,  
Alexandria University, Egypt  
Email: ymustafa@alexu.edu.eg

**Abstract**—Communication and networking research introduces new protocols and standards with an increasing number of researchers relying on real experiments rather than simulations to evaluate the performance of their new protocols. A number of testbeds are currently available for this purpose and a growing number of users are requesting access to those testbeds. This motivates the need for better utilization of the testbeds by allowing concurrent experimentations. In this work, we introduce a novel mapping algorithm that aims to maximize wireless testbed utilization using frequency slicing of the spectrum resources. The mapper employs genetic algorithm to find the best combination of requests that can be served concurrently, after getting all possible mappings of each request via an induced sub-graph isomorphism stage. The proposed mapper is tested on grid testbeds and randomly generated topologies. The solution of our mapper is compared to the optimal one, obtained through a brute-force search, and was able to serve the same number of requests in 82.96% of testing scenarios. Furthermore, we show the effect of the careful design of testbed topology on enhancing the testbed utilization by applying our mapper on a carefully positioned 8-nodes testbed. In addition, our proposed approach for testbed slicing and requests mapping has shown an improved performance in terms of total served requests, about five folds, compared to the simple allocation policy with no slicing.

**Keywords**—wireless testbed, genetic algorithm, frequency slicing, virtualization, induced sub-graph isomorphism

## I. INTRODUCTION

Developing new communication and network protocols and standards is the focus of tremendous number of research studies. To evaluate the performance of these protocols and standards, researchers have been conducting theoretical analysis, simulation, and/or real experimentation on testbeds. Theoretical analysis and simulations suffer in most cases from over simplifications. Moreover, simulations are affected by the complexity of the simulation framework. Therefore, it is more credible for researchers to verify their ideas on real testbeds. Many academic testbeds are currently available for this purpose, such as Orbit [1], GENI [2], Emulab [3]. An increasing number of users are demanding access to those testbeds which necessitates the sharing of testbed resources among multiple concurrent users [4]. While there are many techniques on wired testbeds for resources slicing to share/isolate different experiments efficiently, the problem is more challenging in wireless testbeds. Therefore, there is a need for efficient automated mapping and scheduling algorithms for wireless

testbeds that aim to maximize the number of concurrent non-conflicting slices.

Currently, most of the mapping techniques introduced in the literature target wired testbeds. The problem of assigning nodes to user requests in a wired-connected testbed without violating link constraints is known to be NP-hard [5]. Fair sharing of overall testbed resources (bandwidth, memory, and computational capabilities) among users is a big challenge as well. The main difference between the mapping problem in a wired versus a wireless environment is in the slicing of the testbed. While wired resources can be sliced efficiently, and interfaces can be virtualized with minimal effect, sharing a wireless interface may affect the sharing of interfaces on other nodes because of the wireless resources' inter-dependencies. To guarantee the isolation of experiments, wireless testbeds such as Orbit allocates the entire testbed to one user per time (until now). Other wireless testbeds that enables sharing as NITOS [6], relies on on-demand static allocation based on frequency slicing. This paper presents a novel mapping technique based on genetic algorithm that considers the testbed topology and utilizes spectrum slicing and resource virtualization to map multiple concurrent users maximizing the wireless testbed utilization. The contributions of this paper are: introducing an automated mapping algorithm that benefits from virtualization and spectrum slicing to maximize the number of concurrent users on a wireless testbed, a quantitative evaluation to highlight the advantages and performance of the algorithm, and revealing the possible extensions of our mapper.

The rest of this paper is organized as follows. A review of related work is discussed in Section II. The testbed network model and the tackled problem are described in Section III. The proposed genetic algorithm-based mapper is detailed in Section IV. Evaluation results of the algorithm are shown on Section V. Finally, conclusion and future work are presented in Section VI.

## II. RELATED WORK

The wireless testbed mapping problem is similar to the previous works on Virtual Network Embedding (VNE) on shared wired testbeds [7]–[9]. However, on a wired topology, requests generally consists of bandwidth requirements between node pairs and node capacity, without accounting for interference on a wireless medium or its effect on the maximum number

of allowed virtual machines on a node. Thus, those algorithms mainly map a required virtual topology as source/destination pairs. On the other hand, mapping multiple concurrent users on a wireless testbed differs in how they can share the wireless medium itself.

The standard VNE problem is known to be NP-hard because of its constraints [5]. Many approaches were introduced in the literature to solve this problem. Yu et al. [10] proposed a two-stage solution in which they first map virtual nodes then map virtual links, assuming that path splitting is supported by the underlying network. They then employ path migration to periodically re-optimize the utilization of the substrate network. Chowdhury et al. [8] considered location requirements as well and solved the problem using Mixed Integer Program (MIP) formulation, offering better coordination between the two stages.

Other approaches as vnmFlib network mapping library [9], implemented by Lischka and Karl, used the fact that the network mapping problem resembles the subgraph isomorphism detection, and developed a backtracking algorithm based on the VF2 algorithm [11] used in the pattern recognition community for finding subgraph isomorphisms in large graphs. We inherit this approach to pre-find the possible mappings for each requested topology as a first stage of our mapper algorithm. VnmFlib maps virtual nodes and links in a single stage with better efficiency. VF2x [12] was proposed to enhance vnmFlib and solve some of its limitations.

The Emulab testbed mapping algorithm, Assign [7], considers the online embedding problem, where mapping decisions are taken for each incoming request, with bandwidth constraint alongside constraints on exclusive use of nodes. It categorizes resources into equivalence classes to reduce search space and uses simulated annealing to find a good match. Singhal thesis work [13] targets network virtualization to allow more than one user to share the same wireless node, using a User Mode Linux (UML) network virtualization technique, which he described and evaluated.

In [14], Mahindra et al. provide a comparison between two types of testbed slicing: the usage of flexible time division versus space separation for small sets of nodes, favoring space over time division. The spectrum slicing is introduced in [15] and [16] in the context of wireless standard 802.11, intending to provide each experiment a "spectrum slice" that minimally interferes with other slices, through the use of orthogonal channels. The authors use a spectrum slicing technique on Wi-Fi nodes that can be used on OMF-based testbed [17], targeting Orbit. This spectrum sharing technique is employed for slicing testbed resources in our CRC testbed developed in Alexandria University, Egypt [18].

Although many works target virtualization of the wireless interfaces [19]–[24], most algorithms add limitations to the type of experiments that can be performed, such as delay, overhead, time separation, space separation, transmission power, synchronization, limited topologies, and many others. Due to the difficulties of implementing general solution for interface virtualization, we rely on assigning each of the

wireless interfaces of a node to at most one user as a basic assumption in our testbed nodes modeling. Hence, a node, not interface, virtualization is allowed. In addition, the testbed spectrum resources are shared between users by assigning different wireless channels to distinct users whenever possible. This virtualization scheme maximizes testbed utilization while achieving separation between various experiments with no added limitations. Our work is the first to our knowledge that maps users' requests to enable concurrent experiments on wireless testbeds, assigning their requested topology, interfaces, channels and other preferences with no human intervention. It can be applied on existing grid or non-grid testbeds.

### III. TESTBED MODEL AND PROBLEM DESCRIPTION

The target of our mapper is to maximize the testbed utilization by determining the best subset of requests that can coexist concurrently without interfering with each other. The problem is to find which subset of nodes can serve a user's requested topology, and at the same time, fit with as many other users as possible using the spectrum and hardware resources of the testbed. Furthermore, a priority for each request is included in our utility function that may favor serving one request among others (for example, one belonging to senior researchers group versus students).

The mapper is implemented as two consecutive stages, the first stage is the induced subgraph isomorphism stage which finds all the possible mappings of every requested topology to our testbed. The second stage is a heuristic search using genetic algorithm to find the combination of requests and mappings that lead to the best utilization. Our algorithm is described in more details in the next section.

#### A. Testbed physical model

For our mapping problem, the testbed is viewed as group of  $N$  wireless nodes, each having  $i$  different types of wireless interfaces (for example, if a node has one WiFi card and two types of software defined radio interfaces, then  $i = 3$ ). A testbed reservation request specifies the number of nodes, their specifications, and connectivity. Each node can be either a physical node, with all its resources and network interfaces, or a virtual node with slice of computational and memory resources and one or more of the node's wireless interfaces. The developed reservation system benefits from virtualization to allow multiple concurrent testbed users if possible. A node can be shared between up to  $i$  virtual machines depending on the requests details, and which interfaces are requested.

The topology of the testbed is represented as a square  $N \times N$  binary connectivity matrix, with each element  $(j, k)$  representing the link between nodes  $j, k$ , i.e., whether nodes  $j$  and  $k$  are within the wireless range of each or not. This can be determined empirically through a series of transmission and reception between the testbed nodes to identify the coverage of each node, simplified by defining a threshold for power received for a connected link. Another matrix is used to identify the interfaces that satisfy this connectivity matrix. These two matrices are used to determine which nodes/interfaces can serve a particular request with specific topology and frequencies.

TABLE I: Request's Mathematical Notations

Symbol	Description
$N_r$	The number of requested nodes
$Type_r$	Whether the nodes are physical or virtual
$Int_r$	Requested wireless interfaces
$Ch_r$	Frequency channels requested per interface
$G_r$	The required topology represented by $N_r \times N_r$ connectivity (adjacency) matrix
$TS_r$	Number of requested time-slot(s)
$P_r$	Priority of the request as assigned by the system depending on the user category/priority.

Each of the  $i$  interfaces has a number of orthogonal channels depending on its type, defined by  $max\_ch_i$ , and the channels requested by all requests served in a time-slot cannot exceed that limit to avoid interference. We allow frequency slicing, i.e. sharing the spectrum between different requests when interference can be avoided.

### B. Requests

We have a set  $R$  of  $n$  requests demanding testbed resources in a time-slot, where  $n = |R|$ . A set of parameters are associated with each request  $r$  to identify its details. These parameters are: requested time-slots  $TS_r$ , priority of the request  $P_r$ , requested wireless interfaces  $Int_r$ , requested frequency channels per interface  $Ch_r$ , the number of requested nodes  $N_r$ , whether they are virtual or physical  $Type_r$ , and the required topology represented by another connectivity matrix  $G_r$ . A user can request either specific fixed channels, or a number of flexible channels, for each requested interface. Table I summarizes those notations for use in next sections.

## IV. GA-BASED MAPPER

In this section, the two stages of our proposed mapping algorithm are described.

### A. Generating Candidate Mappings

The first stage of the mapping algorithm is to generate candidate mappings for each requested topology using induced subgraph isomorphism techniques. The induced subgraph isomorphism problem can be defined as follows: given  $H$  and  $G$  graphs, determine whether there is an isomorphic mapping from  $H$  to  $G$ , and find such mappings if exist. The difference of induced from general non-induced subgraph isomorphism problem is that the absence of an edge (link) in  $H$  implies that the analogous edge in  $G$  is also absent. In the non-induced problem those edges in  $G$  can be present. The induced subgraph isomorphism fits our needs in wireless testbeds as we want to provide the requested links of the user topology and avoid potential interference of any other links.

We use igraph [25] library to get all isomorphic induced sub-graphs of each request. The subgraph  $H$  represents the requested topology( $G_r$ ), while  $G$  denotes the testbed connectivity. After this step, several mappings are available - if exist - associated with each  $G_r$ . Those distinct mappings are used to maximize the odds of coexistence between different requests.

TABLE II: Chromosome Structure

$r_1$	$r_2$	...	$r_n$
1	0	...	1
3	4	...	6

### B. Genetic Algorithm

The second stage of the mapper is the genetic algorithm-based stage (GA-based stage). After getting all the candidate mappings of each request, we apply genetic algorithm for choosing the subset of requests that can run together without conflicts while satisfying the constraints. The details of the GA-based stage of the mapper are described in the following subsections.

1) *Chromosome Structure*: A GA chromosome consists of  $n$  genes, each representing a particular request, where  $n$  is the total number of requests. Each gene consists of two fields:

- A bit representing whether the corresponding request shall be allocated or not in the current solution.
- An integer representing a pointer to a mapping chosen from the candidate mappings of this request as determined from the first stage.

For example, in Table II,  $r_1$  and  $r_n$  will both be served with their 3rd and 6th mapping respectively, while  $r_2$  with its 4th mapping will not be served.

2) *Initialization*: First, we generate the population of chromosomes with random values. Each gene (request) in the chromosome is initialized randomly for both its existence and mapping. The same gene on different chromosomes can be linked with different mapping of its candidate mappings (if there are many).

Then, the requested frequencies by all requests are checked for any conflicts. If no conflict is indicated for all requests, then this step is not needed to be done for any subgroup of the requests later during the algorithm. If a conflict is indicated, then only subset of the requests can be served together and this check has to be performed for each chromosome (subset of requests) in every iteration of the GA.

3) *Selection*: The fitness function in eq. 1 reflects the priority of the request, the requested time duration, and the penalties resulting from any conflicts. In our case, the genetic algorithm tries to find the chromosome which minimizes that fitness function. Two types of penalties are calculated: resource penalties and channel penalties. If multiple requests ask for the same resource (node or interface on a particular node), the number of resource conflicts is increased. Moreover, if multiple requests demand the same frequency channel, the number of channel conflicts is increased. The penalties term is the multiplication of these conflicts by a large number.

The fitness function is,

$$Fitness = Penalties + \sum_{r=1}^n \left( \frac{w_1}{P_r} + w_2 \times TS_r \right) \quad (1)$$

$$Penalties = \sum_{conflicts} (Resource\_conflicts + Channel\_conflicts) \times Large\_number \quad (2)$$

where  $w_1$  and  $w_2$  are adjustable parameters used to indicate the relative importance of a term over another. In order to understand this formulation, let us consider the case when all requests have equal priorities, have the same time preferences and no penalties occur. In this case, the problem turns to optimization problem that aims to maximizing the number of requests being served on this slot. The fitness function also allow us to serve a request which consumes larger resources or longer duration if it has higher priority than other request(s).

4) *Genetic Operations*: The two primary genetic operations, mutation and crossover, are applied to chromosomes randomly with probabilities  $p_m$  and  $p_c$  respectively.

- Mutation: Two types of mutations are applied, each with probability 0.5:
  - Mutation 1: randomly toggle the first field boolean of a gene from 0 to 1 and vice versa keeping the same mapping.
  - Mutation 2: randomly change the mapping of a request to another one of its candidate mappings.
- Crossover: Apply a single-point crossover between chromosome pairs.

5) *Termination*: When a maximum number of generations is reached, or no significant relative change in the fitness over a predefined stall generations happens, the algorithm is terminated and the best solution that corresponds to the best chromosome of the last generation is determined.

## V. EVALUATION

In the current section, the simulation framework and the different scenarios used to evaluate the performance of the proposed mapper are described and the results are presented.

### A. Simulation Framework

In order to evaluate the performance of our GA-based mapper, we setup a simulation environment for different testbed sizes and parameters. We compare our mapping algorithm to the optimal solution found using the exhaustive search. For this purpose, a brute-force search algorithm (BF) was implemented to find the best possible solution, optimum, given a particular subset of requests. The algorithms are evaluated and compared under multiple testbed sizes and mappings size limits. Moreover, the revenue of applying our GA-based mapper with enabled spectrum slicing and resources virtualization over a simple allocation policy with no slicing is demonstrated while using all possible mappings provided by the induced graph isomorphism algorithm.

The performance metrics considered are:

- i) *Number of served requests*, the average number of accepted requests by the mapper to be served.
- ii) *Optimality*, the percentage of solutions obtained by our mapper that were identical to those found through brute-force search for the optimal solutions.
- iii) *Slicing revenue*, the increased testbed utilization, in terms of the number of concurrent users that our mapper can allocate relative to the simple allocation policy with no slicing (one user).

TABLE III: Values of the parameters used for evaluation

Parameter	Distribution/Value
$pop - size$	60
$p_m$	0.2
$p_c$	0.8
$max - iter$	500
$i$	3 interfaces
$Max\_ch_i$	13 for $i = 1$ , 13 for $i = 2$ , and 40 for $i = 3$ .
$Ch_r$	$\sim \mathcal{N}(\mu, \sigma^2)$ where $\mu = 0.25 \times Max\_ch_i$ , and $\sigma = 1/6 \times Max\_ch_i$ . The distribution is limited to half of $Max\_ch_i$ .
$P_r$	$\sim U(1, 5)$ .
$TS_r$	$\sim \mathcal{N}(\mu, \sigma^2)$ where $\mu = 2$ , and $\sigma = 5$ .
$Int_r$	$\sim U(1, 3)$

### B. Experimental Set-up


To evaluate our mapping algorithm, random input data are generated to represent several users with variable request preferences. This subsection describes the experimental setup and the statistics of the parameters used.

A wireless testbed physical topology is expressed as graph with nodes represent the testbed nodes and links represent the connectivity of nodes. This information is embedded as a connectivity/adjacency matrix. For the experiments, testbed nodes were assumed to assemble either a grid topology or an arbitrary randomly generated topology. In grid testbeds, each node can only see one neighbor in each direction, and in random topology, the number of connections are random. Different testbed sizes have been used of  $3 \times 3$  grid (9 nodes),  $4 \times 4$  grid (16 nodes),  $5 \times 5$  grid (25 nodes), and  $6 \times 6$  grid (36 nodes).


The number of nodes of each request  $N_r$  were picked from a uniform distribution with mean roughly equals to half of the available nodes on smallest testbed (between three and five nodes). For each request, we generate a random adjacency matrix  $G_r$  representing its requested topology. This range of number of nodes per request is selected as when the number of nodes exceeds five, most of the randomly generated topologies does not fit on the testbed as indicated from the sub-graph isomorphism stage.

The requests parameters  $Ch_r$ ,  $P_r$ , and  $TS_r$  where randomly generated for each request to exhibit different users' categories. The number of requested channels  $Ch_r$  were selected from a normal distribution limited to half of the maximum number of channels available at most. Five different categories of users were assumed, and hence five uniformly distributed priorities. Moreover, each request demands a random number of time-slots distributed normally around two slots with a high standard deviation of five. The population size of the GA is 60 and the crossover rate  $p_c$  and mutations rate  $p_m$  were set to 0.8 and 0.2 respectively. This relatively high mutation rate was used to avoid being trapped in local optima. The GA is terminated after 500 iterations (no further improvement has been found for using more iterations). Table III summarizes the parameters values and distributions used for evaluation.


For the first scenario, comparison between our mapper solution and the optimal obtained by brute force search, the number of




(a) Physical requests, 10 Mappings.



(b) Virtual requests enabled, 10 Mappings.



(c) Physical requests, 15 Mappings.



(d) Virtual requests enabled, 15 Mappings.

Fig. 1:  $1^{st}$  scenario I: Average number of served requests for grid topologies. Mapping limits of 10 and 15, results, for requests of both virtualization enabled and of only physical nodes.

mappings of a topology  $G_r$  produced by the first stage of the mapper will be limited. We tried limiting the solutions to ten and fifteen mappings only, for each incoming request. In addition, the number of users demanding testbed resources  $|R|$  is fixed to five. These were done to shrink the search space as increasing it would require enormous computation resources and large amount of time for a single run of the brute-force implementation. Two settings were tested: first, with no virtualization enabled, i.e.  $Type = 0$ . Second, enabling requests demanding virtual nodes,  $Type = 1$ . Each experiment is repeated 10 times, with different random inputs, and the average value and variance of these repetitions are presented in the simulation results of the following figures.

Then, we do the simulations for the second scenario, considering all possible mappings for each incoming request and running the GA-based mapper. The effect of varying testbed sizes will be shown, as well as varying the number of requests on the slicing revenue. Similarly, each experiment is repeated 10 times, and the average value of these repetitions is presented in the simulation results. Furthermore, the execution time of both GA and BF algorithms are depicted.

### C. Performance results

Fig.1 shows the average number of requests to be served for different mapping limits of 10 and 15, for both  $Type = 1$  (virtualization-enabled) and  $Type = 0$  (only physical) on the grid testbeds. In addition, Fig.2 presents the results for randomly generated topologies of the same-sizes testbeds. We notice three bars on Fig.1(a), 1(c) and Fig.2(c), with no error bars, where the results over ten runs exhibits no variations. As seen from the sub-figures, our mapper (GA) serves almost exactly the same number of requests compared






Fig. 3: Percentage of optimal solutions found by GA-mapper on grid and random topologies.

to the brute force (BF) solution. In total, for about **82.96%** of runs, the GA-based mapper serves the same number of requests as the BF. Besides, it is noticeable that the average number of served requests is not strictly increasing with the increase of testbed sizes, which can be attributed to the grid topology limited network connectivity as well as the high average number of requested channels per user. Furthermore, the number of successfully allocated requests increase with the mapping limits, which is intuitive. The more possible different mappings that are fed to the system, the more probable it could find a better solution and allocate more users concurrently. Moreover, the average number of allocated requests on the physical case is more than one, which demonstrates the sharing benefits of the frequency slicing technique used.


To demonstrate the effect of the testbed physical topology on the results of the average number of allocated requests, the same procedure is performed on carefully planned topology similar to the one used by the CRC testbed [18]. The nodes graph of this 8-nodes network is shown on Fig.4. The results obtained on this topology, with five requests with the same




(a) Physical requests, 10 Mappings.



(b) Virtual requests enabled, 10 Mappings.



(c) Physical requests, 15 Mappings.



(d) Virtual requests enabled, 15 Mappings.

Fig. 2: 1<sup>st</sup> scenario II: Average number of served requests for random topologies. Mapping limits of 10 and 15, results, for requests of both virtualization enabled and of only physical nodes.




Fig. 4: The topology of the adopted 8-nodes wireless testbed.

parameters, and mapping limits of 10 and 15 are displayed on Fig.5. As observed from the figure, applying our mapper on this topology results in increase in number of served requests for a smaller sized testbed. This highlights the fact that the grid topology is not the best way to organize testbed nodes. Fig.3 shows the optimality metric results for both grid and random topologies testbeds with variable number of nodes in the testbed. The figure indicates a good steady performance with increased testbed sizes, because of the careful design of chromosomes embedding we used, with an average of **65%** over all runs.

We manifest the slicing revenue resulted from resources slicing and our mapping algorithm on the testbed utilization on Fig.6 and Fig.7, using all mappings of the first stage of our approach. The figures show the effect of varying the incoming number of requests on the revenue, and the impact of varying the testbed size as well for both grid and random topologies. Our adopted frequency-based slicing scheme considered by our mapper offers up to 490% increase in the number of concurrent




Fig. 5: Average number of served requests for mapping limits of 10 and 15 for the 8-node testbed, virtualization enabled.

users relative to the simple allocation scheme where no slicing supported on grid topologies, and up to 500% increase on random topologies.

Finally and as expected, our algorithm solves this optimization problem pretty fast relative to the exhaustive (BF) search algorithm. As depicted in Fig.8, the exhaustive search execution time increases exponentially with the increase in search space (mappings limit) while our algorithm scales almost linearly.

## VI. CONCLUSION AND FUTURE WORK

This paper presents our implementation of a mapping algorithm for reservation requests on wireless testbeds, based on induced sub-graph isomorphism and genetic algorithm. We describe the concept and illustrate its performance using simulation for different scenarios. In almost all of the experiments performed, our genetic algorithm-based mapper *succeeded to allocate the same number of requests allocated by the best solution found by brute force search in 82.96% cases*. We manifested that a careful choice of nodes locations of the testbed can enhance the number of users that can be concurrently served. Moreover, our proposed technique




Fig. 6: Slicing Revenue for different number of requests and different grid testbed sizes.




Fig. 7: Slicing Revenue for different number of requests and different testbed sizes of random topologies.

achieves up to 500% revenue of testbed users by the suggested slicing mechanism compared to no-slicing strategy.

We are working on integrating a 24 hours scheduler into the current implementation of mapping algorithm. The fitness function will be re-designed to include scheduling details, therefore, performing both mapping and providing the daily schedule of requests for each time-slot.

#### REFERENCES

- [1] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols," in *IEEE Wireless Communications and Networking Conference*, vol. 3. IEEE, 2005, pp. 1664–1669.
- [2] M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, R. Ricci, and I. Seskar, "Geni: A federated testbed for innovative network experiments," *Computer Networks*, vol. 61, pp. 5–23, 2014.
- [3] "UTAH Emulab website," <http://emulab.net>.
- [4] D. Stavropoulos, A. Dadoukis, T. Rakotoarivelo, M. Ott, T. Korakis, and L. Tassioulas, "Design, architecture and implementation of a resource discovery, reservation and provisioning framework for testbeds." WiN-MeE, 2015.
- [5] D. G. Andersen, "Theoretical approaches to node assignment," *Computer Science Department*, p. 86, 2002.
- [6] "UTH NITOS website," <http://nitlab.inf.uth.gr/NITlab/index.php/testbed>.
- [7] R. Ricci, C. Alfeld, and J. Lepreau, "A solver for the network testbed mapping problem," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 2, pp. 65–81, 2003.
- [8] N. Chowdhury, M. R. Rahman, and R. Boutaba, "Virtual network embedding with coordinated node and link mapping," in *IEEE INFOCOM 2009*. IEEE, 2009, pp. 783–791.
- [9] J. Lischka and H. Karl, "A virtual network mapping algorithm based on subgraph isomorphism detection," in *Proceedings of the 1st ACM workshop on Virtualized infrastructure systems and architectures*. ACM, 2009, pp. 81–88.
- [10] M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking virtual network embedding: substrate support for path splitting and migration," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 17–29, 2008.




Fig. 8: Elapsed execution time of brute-force versus GA-based mapper for different mapping limits.

- [11] L. P. Cordella, P. Foggia, C. Sansone, and M. Vento, "A (sub) graph isomorphism algorithm for matching large graphs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 10, pp. 1367–1372, 2004.
- [12] Q. Yin and T. Roscoe, "Vf2x: fast, efficient virtual network mapping for real testbed workloads," in *Testbeds and Research Infrastructure Development of Networks and Communities*. Springer, 2012, pp. 271–286.
- [13] S. Singhal, G. Hadjichristofi, I. Seskar, and D. Raychaudhuri, *Evaluation of UML based wireless network virtualization*. IEEE, 2008.
- [14] R. Mahindra, G. Bhanage, G. Hadjichristofi, I. Seskar, D. Raychaudhuri, and Y. Zhang, "Space versus time separation for wireless virtualization on an indoor grid," in *Next Generation Internet Networks, 2008. NGI 2008*. IEEE, 2008, pp. 215–222.
- [15] A.-C. Anadiotis, A. Apostolaras, D. Syrivelis, T. Korakis, L. Tassioulas, L. Rodriguez, and M. Ott, "A new slicing scheme for efficient use of wireless testbeds," in *Proceedings of the 4th ACM international workshop on Experimental evaluation and characterization*. ACM, 2009, pp. 83–84.
- [16] A.-C. Anadiotis, A. Apostolaras, D. Syrivelis, T. Korakis, L. Tassioulas, L. Rodriguez, I. Seskar, and M. Ott, "Towards maximizing wireless testbed utilization using spectrum slicing," in *Testbeds and Research Infrastructures. Development of Networks and Communities*. Springer, 2011, pp. 299–314.
- [17] T. Rakotoarivelo, M. Ott, G. Jourjon, and I. Seskar, "Omf: a control and management framework for networking testbeds," *ACM SIGOPS Operating Systems Review*, vol. 43, no. 4, pp. 54–59, 2010.
- [18] "CRC testbed website," <http://smartci.alexu.edu.eg/crc/doku.php>.
- [19] G. Bhanage, D. Vete, I. Seskar, and D. Raychaudhuri, "Splitap: leveraging wireless network virtualization for flexible sharing of wlangs," in *IEEE GLOBECOM*. IEEE, 2010, pp. 1–6.
- [20] S. S. Hong, J. Mehlman, and S. Katti, "Picasso: flexible rf and spectrum slicing," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 37–48, 2012.
- [21] L. Wang, Z. Zhang, L. Ye, B. Gong, and W. Lou, "Fully reconfigurable cognitive ofdma testbed," in *Communications and Networking in China (CHINACOM), 2011 6th International ICST Conference on*. IEEE, 2011, pp. 933–938.
- [22] A. J. Nicholson, S. Wolchok, and B. D. Noble, "Juggler: Virtual networks for fun and profit," *IEEE Transactions on Mobile Computing*, vol. 9, no. 1, pp. 31–43, 2010.
- [23] A. D. Rivera and W. Zucci, "Virtualization of wireless network interfaces wi-fi ieee 802.11," in *Proceedings of the 9th WSEAS International Conference on Telecommunications and Informatics*, 2010, pp. 46–51.
- [24] C. Liang and F. R. Yu, "Wireless network virtualization: A survey, some research issues and challenges," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 358–380, 2015.
- [25] "igraph - the network analysis package," <http://igraph.org>.